OIPE

OCT 2 3 2003

PATENT & TRADEMARK OFFICE

AF/2700

TRANSMITTAL OF APPEAL BRIEF (Large Entity)

Docket No.
POU920000009US1

In Re Application Of: **Novaes et al.**

Serial No.	Filing Date	Examiner	Group Art Unit
09/583,784	05/31/2000	Hassan Mahmoudi	2175

Invention: **METHOD, SYSTEM AND PROGRAM PRODUCTS FOR RECOVERING FROM FAILURES WITHIN A SHARED NOTHING DISTRIBUTED COMPUTING ENVIRONMENT**

RECEIVED

OCT 2 8 2003

Technology Center 2100

TO THE COMMISSIONER FOR PATENTS:

Transmitted herewith in triplicate is the Appeal Brief in this application, with respect to the Notice of Appeal filed on

The fee for filing this Appeal Brief is: **$330.00**

☐ A check in the amount of the fee is enclosed.

☐ The Director has already been authorized to charge fees in this application to a Deposit Account.

☒ The Director is hereby authorized to charge any fees which may be required, or credit any overpayment to Deposit Account No. **09-0463 (IBM)**

Kevin P. Radigan

Signature

Dated: **October 21, 2003**

Kevin P. Radigan, Esq.
Registration No.: 31,789
Heslin Rothenberg Farley & Mesiti P.C.
5 Columbia Circle
Albany, New York 12203
Telephone: (518) 452-5600
Facsimile: (518) 452-5579

I certify that this document and fee is being deposited on **October 21, 2003** with the U.S. Postal Service as first class mail under 37 C.F.R. 1.8 and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Kevin P. Radigan

Signature of Person Mailing Correspondence

Kevin P. Radigan, Esq.

Typed or Printed Name of Person Mailing Correspondence

cc:

#15
11/3/03
A.W.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Appellants:	Novaes et al.	:	Group Art Unit: 2175
Serial No.:	09/583,784	:	Examiner: Hassan Mahmoudi
Filed:	05/31/2000	:	Appeal No.:

Title: METHOD, SYSTEM AND PROGRAM PRODUCTS FOR RECOVERING FROM FAILURES WITHIN A SHARED NOTHING DISTRIBUTED COMPUTING ENVIRONMENT

Certificate of Mailing

I hereby certify that this correspondence is being deposited with the U.S. Postal Service as first class mail in an envelope addressed to: Mail Stop Appeal Briefs – Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on October 21, 2003.

Kevin P. Radigan

Kevin P. Radigan
Attorney for Appellants
Registration No.: 31,789

Date of Signature: October 21, 2003.

Mail Stop Appeal Briefs – Patents
Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Brief of Appellants

Dear Sir:

This is an appeal from a final rejection, dated May 22, 2003, rejecting claims 1-25, all the claims being considered in the above-identified application. This Brief is accompanied by a transmittal letter authorizing the charging of appellants' deposit account for payment of the requisite fee set forth in 37 C.F.R. §1.17(c).

Real Party in Interest

This application is assigned to **International Business Machines Corporation** by virtue of an assignment executed by the co-inventors on September 7, 2000 and September 26, 2000, and recorded with the United States Patent and Trademark Office at reel 011186, frame 0731, on October 10, 2000. Therefore, the real party in interest is **International Business Machines Corporation**.

Related Appeals and Interferences

To the knowledge of the appellants, appellants' undersigned legal representative, and the assignee, there are no other appeals or interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the instant appeal.

Status of Claims

This patent application was filed on May 31, 2000 with the United States Patent and Trademark Office. As filed, the application included three (3) independent claims (i.e., claims 1, 2 & 3). On April 12, 2002, appellants filed a Preliminary Amendment that added (dependent) claims 4-24.

In an initial Office Action dated May 8, 2002, claims 1-24 were rejected under 35 U.S.C. 103(a) as being unpatentable over Badovinatz et al. (U.S. Patent No. 5,805,786) in view of Tsukerman et al. (U.S. Patent 6,341,340), and further in view of Cotner et al. (U.S. Patent 5,884,327). In appellants' response dated September 3, 2002 (with which a request for a one-month extension of time and requisite fee were enclosed), claims 1, 2 & 3 were amended.

In a second Office Action dated December 4, 2002, claims 1-6, 11-13 & 18-20 were rejected under 35 U.S.C. 103(a) as being unpatentable over Dias et al. (U.S. Patent No. 5,907,849) in view of Blott et al. (U.S. Patent No. 6,449,618); and claims 7-10, 14-17 & 21-24 were rejected under 35 U.S.C. 103(a) as being unpatentable over Dias et al. in view of Blott et al.

as applied to claims 1-6, 11-13 & 18-20, and further in view of Badovinatz et al. In appellants' response dated February 28, 2003, claim 25 was added.

In a third and final Office Action dated May 22, 2003, claims 1-6, 11-13 & 18-20 were rejected under 35 U.S.C. 103(a) as being unpatentable over Dias et al. in view of Blott et al.; claims 7-10, 14-17 & 21-24 were rejected under 35 U.S.C. 103(a) as being unpatentable over Dias et al. in view of Blott et al., as applied to claims 1-6, 11-13 & 18-20, and further in view of Badovinatz et al.; and claim 25 was rejected under 35 U.S.C. 103(a) as being unpatentable over Dias et al. in view of Blott et al., as applied to claims 1-6, 11-13 & 18-20, and further in view of Carter et al. (U.S. Patent No 5,909,540). In appellants' response dated July 22, 2003, no claims were amended.

A Notice of Appeal to the Board of Patent Appeals and Interferences was filed on August 21, 2003. The status of the claims is therefore as follows:

Claims allowed – none;

Claims objected to – none;

Claims rejected – 1-25; and

Claims canceled – none.

Appellants are appealing the rejection of claims 1-25.

Status of Amendments

Pursuant to the Advisory Action dated August 11, 2003, appellants' remarks proffered in the Response to the final Office Action dated July 22, 2003 were entered upon filing of the Notice of Appeal and this Appeal Brief. However, no claim amendment was effectuated by the Response. The claims as set out in Appendix A include all prior entered amendments.

Summary of Invention

The present invention is directed to the problem of recovering from failures at nodes within a distributed computing environment.

More particularly, the present invention is a method (claim 1), system (claim 2), and program storage device, embodying a program of instructions readable and executable by a machine to perform the method (claim 3), for recovery from failures within a shared nothing distributed computing environment (see page 35, lines 21-25). The method includes detecting a failure 1500 (FIG. 15; page 20, lines 25-28) within the shared nothing distributed computing environment; and automatically recovering from the failure, wherein one or more transactions affected by the failure are executed to completion without rolling back the one or more transactions (see page 32, lines 20-23; FIG. 26; and page 33 line 7 to page 35 line 5) and without requiring a reposting of the one or more transactions (page 32, lines 10-19).

In another aspect of the invention (see claim 7), wherein the shared nothing distributed computing environment comprises a processing group 500 (FIG. 5; see page 12, lines 9-18) with a plurality of members 502 (FIG. 5), failure recovery includes electing a coordinator from among the (at least one) surviving member(s) (see page 33, lines 16-17). The coordinator receives messages pertaining to one or more transactions from the other surviving members 2600 (FIG. 26; page 33, lines 21-25). (See claims 8 and 9.)

In a further aspect (see claim 25), the shared nothing distributed computing environment comprises a distributed synchronous transaction system (see page 10, lines 8-12 and 17-24; 204, FIG. 2), and the method includes a failure recovery method for the distributed synchronous transaction system (see page 32, lines 20-25).

Issues

1. Whether claims 1-6, 11-13 & 18-20 were rendered obvious under 35 U.S.C. 103(a) by Dias et al. in view of Blott et al.

2. Whether claims 7-10, 14-17 & 21-24 were rendered obvious under 35 U.S.C. 103(a) by Dias et al. in view of Blott et al., as applied to claims 1-6, 11-13 & 18-20, and further in view of Badovinatz et al.

3. Whether claim 25 was rendered obvious under 35 U.S.C. 103(a) by Dias et al. in view of Blott et al., as applied to claims 1-6, 11-13 & 18-20, and further in view of Carter et al.

Grouping of Claims

As to the rejections applied against claims 1-25, it is appellants' intention that the rejected claims do not stand or fall together. For example, since each ground of rejection provides a grouping of claims, the following principal groups of claims are included herein:

I. Claims 1-6, 11-13 & 18-20;

II. Claims 7-10, 14-17 & 21-24; and

III. Claim 25.

As understood, the claims of one group of claims do not stand or fall with any other group of claims. Further, appellants' respectfully submit that the claims within each group do not stand or fall together.

Argument

Group I: Claims 1-6, 11-13 & 18-20

As noted, claims 1-6, 11-13 & 28-20 stand rejected under 35 U.S.C. §103(a) as being obvious over Dias et al. in view of Blott et al. Appellants respectfully request reversal of this rejection.

Absent from the final Office Action is an express teaching, suggestion or incentive identified in the art for making the combination. The justification for the combination of patents in the final Office Action is that they allegedly teach the benefit of making the recovery process

faster and more efficient by avoiding manual intervention, roll back or reposting of the recovered transactions. Appellants respectfully submit that this justification does not identify an adequate teaching, suggestion or incentive in the art itself for the combination proposed in the final Office Action. In this case, the basis for the combination is believed drawn from appellants' own disclosure. Appellants' above-summarized technique comprises an approach for recovery from failures within a shared nothing distributed computing environment wherein one or more transactions affected by a failure are automatically executed to completion. It is appellants' disclosure which teaches that this execution is accomplished without rolling back and without reposting the one or more transactions. The protocol necessary to accomplish such a recovery and automatic execution of that transaction would not be readily apparent to one skilled in the art given the teachings of Dias et al. and Blott et al.

Dias et al. describe a technique for recovering from a failure of a processing node in a partitioned shared nothing database processing system. The system described by Dias et al. is a high availability system, meaning that there cannot be a single point of failure. For example, reference column 1, lines 39-43 thereof. In contrast, the teachings of Blott et al. are directed to a real-time event processing system for processing a sequence of events generated by one or more applications. The Blott et al. system is an architecture which has a single point of failure on which the whole system depends. For example, reference column 8, lines 6-8. Thus, appellants respectfully submit that Blott et al. is deficient in teaching any recovery approach for a high availability system, such as described by Dias et al. Because the only approach described by Blott et al. comprises a single point of failure approach, appellants respectfully submit that one of ordinary skill in the art would not have been led by the teachings thereof to combine those teachings with Dias et al. as proposed in the Office Action.

For the above reasons, appellants respectfully request reversal of the obviousness rejection to the independent claims presented herewith based upon a proposed combination of Dias et al. and Blott et al.

Even assuming, arguendo, that the combination of Dias et al. and Blott et al. is proper, the combination still fails to teach or suggest features of the recited invention. For example, each of the independent claims presented recites automatically recovering from the failure within the shared nothing distributed computing environment without rolling back the one or more transactions and without requiring reposting of the one or more transactions. These characterizations are believed to distinguish appellants' invention from any combination of Dias et al. and Blott et al.

As indicated in the final Office Action, Dias et al. does not teach automatically recovering from failure, wherein one or more transactions affected by the failure are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions. Appellants agree. For an alleged teaching of this concept, the final Office Action relies upon Blott et al. However, appellants respectfully submit that the final Office Action misinterprets the teachings of Blott et al. in this regard. Blott et al. expressly teach that in nearly all cases it is possible to recover to the last successfully processed event and to replay accepted events from that point forward. (See column 27, lines 23-31). This replaying of accepted events is a reposting of the accepted events, which is clearly contrary to appellants' technique recited in the independent claims presented. Again, appellants recite automatically recovering from the failure both without rolling back the one or more transactions and without requiring a reposting of the one or more transactions. Blott et al. expressly teach that the technique presented replays accepted events from that point forward. This replaying is a reposting of the events, and therefore teaches away from appellants' approach.

In the final Office Action, pages 9 & 10, the Examiner states that Blott et al. teach "filtering duplicate requests" at column 26, lines 6-8, which implies that recovery is done without "reposting" of the transaction. This characterization of the teachings of Blott et al. is respectfully traversed. Blott et al. describe their recovery approach beginning at the bottom of column 26 and continuing onto column 27. As explained in this material, the Blott et al. recovery approach makes use of a feature known as recovery points, which may be viewed as markers that are inserted into a recovery log. A recovery point represents a transactionally consistent state of the

database that can be reverted to as long as the logs are available (see column 27, lines 14-16). Blott et al. expressly teach a logging function wherein events are persistently stored, which enables recovery to the last successfully processed event, and replaying of accepted events from that point forward (see column 27, lines 23-31). This replaying is in fact a reposting of the events, and therefore teaches away from appellants' approach.

The Examiner's citation to column 26, lines 6-8 is believed to be an error with respect to this aspect of appellants' claimed invention. This material in Blott et al. is discussing a backup maintenance procedure or service plan that is to be introduced after each established recovery point. Appellants respectfully submit "recovery" in this instance does not refer to recovery from failure as expressly recited in their independent claims. The daily maintenance and service plan changes discussed at the top of column 26 do not apply to recovery from failure per se, which is expressly discussed in column 27 as described above.

Still further, Blott et al. does not present any recovery technique for a shared nothing distributed computing environment such as recited by appellants. Blott et al. expressly teach the existence of a log that is always accessible to all processing nodes. This assumption does not apply in the case of a shared nothing distributed computing environment. Without this assumption, the recovery approach of Blott et al. would simply be inoperable, and thus, inapplicable, to such a shared nothing environment.

For these additional reasons, appellants respectfully submit that one of ordinary skill in the art would not have been led by the teachings of Dias et al. in combination with Blott et al. to a recovery approach as recited in the independent claims presented. A whole new protocol is added by the present invention for a shared nothing computing environment to accomplish automatic recovery and automatic execution of one or more transactions affected by a failure. This automatic execution is recited to be accomplished without rolling back the transactions and without requiring reposting of the one or more transactions. No similar functionality is provided in either Blott et al. or Dias et al., or the other known patents. Further, appellants respectfully submit that accomplishing the necessary modifications to the approach of Dias et al. and Blott et

al. would not have been readily achievable by one of ordinary skill in the art, as evidenced by the extensive disclosure and drawings of the present application.

For all the above reasons, appellants respectfully request reversal of the rejection of claims 1-6, 11-13 & 28-20.

Group II: Claims 7-10, 14-17 & 21-24

Claims 7-10, 14-17 & 21-24 stand rejected under 35 U.S.C. §103(a) as being obvious over Dias et al. in view of Blott et al., as applied to claims 1-6, 11-13 and 18-20, and further in view of Badovinatz et al. Appellants respectfully request reversal of this rejection.

These dependent claims are believed allowable for the same reasons as the independent claims of Group I, as well as for their own additional characterizations. In this regard, appellants respectfully traverse the secondary obviousness rejections to claims 7-10, 14-17 & 21-24 based on Dias et al. in view of Blott et al., and Badovinatz et al. for the same reasons noted above with respect to the claims of Group I. Briefly summarized, these are: (1) the protocol necessary to obtain automatic failure recovery with transaction completion without rolling back transactions and without reposting transactions affected by the failure would not be readily apparent to one with ordinary skill in the art; and (2) Blott et al. teach replaying of accepted events, which is reposting of the accepted events, so that the proposed combination of Dias et al. and Blott et al. does not teach or suggest automatically recovering from failure without rolling back one or more transactions and without requiring reposting of one or more transactions.

The dependent claims of Group II further recite a processing group with a plurality of members wherein at least one member survives a failure and failure recovery comprises electing a coordinator from among the surviving members of the processing group. Badovinatz et al. was cited in the final Office Action for allegedly teaching a system for managing membership of processor groups in a distributed computing environment (see Col. 1, lines 58-61), including the selection of a new processor as group leader in the event that the current group leader processor fails. (see Col. 5, lines 42-48). Notwithstanding this alleged teaching, appellants respectfully

submit that Badovinatz et al. do not teach the aforementioned deficiencies of Dias et al. and Blott et al. when applied against the claims at issue.

For the reasons stated above, appellants respectfully submit that the claims of Group II would not have been obvious to one with ordinary skill in the art.

Group III: Claim 25

Claim 25 stands rejected under U.S.C. §103(a) as being obvious over Dias et al. in view of Blott et al., as applied to claims 1-6, 11-13 and 18-20, and further in view of Carter et al. Appellants respectfully request reversal of this rejection.

With respect to the rejection of claim 25, appellants respectfully traverse the applicability of the teachings of Carter et al. Carter et al. describe an approach for providing highly available data storage using globally addressable memory. At column 24, lines 34-39, Carter et al. describe the existence of a shared memory subsystem which includes a coherent replication controller. Thus, Carter et al. expressly teach reliance on a shared memory subsystem, which teaches away from applicability to a shared nothing distributed computing environment. In a shared nothing distributed processing system, a shared memory subsystem for the processing nodes is not possible. Thus, appellants respectfully submit that one of ordinary skill in the art would not have extrapolated in some manner the teachings of Carter et al. for combination with a shared nothing distributed computing environment such as recited by appellants.

For all of the reasons stated above, appellants also respectfully submit that claim 25 would not have been obvious to one of ordinary skill in the art based upon the applied art.

Conclusion

Appellants respectfully submit that their claimed invention would not have been obvious to one of ordinary skill in the art based upon Dias et al. and Blott et al., either alone or in

combination, nor would it have been obvious to one of ordinary skill in the art based on these and further in view of either Badovinatz et al. or Carter et al.

By way of example, appellants submit that the combination of Dias et al. and Blott et al. is not proper because Dias et al. describe a high availability system, meaning that there cannot be a single point of failure, while Blott et al. describe a system architecture which has a single point of failure. As another example, Dias et al., Blott et al., and Badovinatz et al., either alone or in combination, do not teach automatically recovering from failures within a shared nothing distributed computing environment, wherein one or more transactions affected by the failure are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions. As still another example, the protocol necessary to obtain automatic failure recovery with transaction completion without rolling back transactions and without reposting transactions affected by the failure would not have been readily apparent to one with ordinary skill in the art based on the applied patents.

Kevin P. Radigan

Kevin P. Radigan
Attorney for Appellants
Registration No.: 31,789

Dated: October 21, 2003.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579

Appendix A

1. A method of recovery from failures within a shared nothing distributed computing environment, said method comprising:

detecting a failure within said shared nothing distributed computing environment; and

automatically recovering from said failure, wherein one or more transactions affected by said failure are automatically executed to completion without rolling back said one or more transactions and without requiring a reposting of said one or more transactions.

2. A system of recovery from failures within a shared nothing distributed computing environment, said system comprising:

means for detecting a failure within said shared nothing distributed computing environment; and

means for automatically recovering from said failure, wherein one or more transactions affected by said failure are automatically executed to completion without rolling back said one or more transactions and without requiring a reposting of said one or more transactions.

3. At least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform a method of recovery from failures within a shared nothing distributed computing environment, said method comprising:

detecting a failure within said shared nothing distributed computing environment; and

automatically recovering from said failure, wherein one or more transactions affected by said failure are automatically executed to completion without rolling back said one or more transactions and without requiring a reposting of said one or more transactions.

4. The method of claim 1, wherein the shared nothing distributed computing environment comprises a processing group with a plurality of members, and wherein the detecting comprises detecting a failure of at least one of the plurality of members.

5. The method of claim 4, wherein the recovering comprises synchronizing messages regarding the one or more transactions among surviving members of the processing group.

6. The method of claim 5, wherein the recovering further comprises committing the one or more transactions.

7. The method of claim 4, wherein at least one member of the processing group survives the failure, and wherein the recovering comprises electing a coordinator from among the at least one surviving member.

8. The method of claim 7, wherein the recovering further comprises receiving by the coordinator a list of the one or more transactions from the other surviving members.

9. The method of claim 8, wherein the recovering further comprises receiving by the coordinator any commit protocol messages for the one or more transactions the coordinator does not already have.

10. The method of claim 9, wherein the coordinator initiates the commit protocol for the one or more transactions.

11. The system of claim 2, wherein the shared nothing distributed computing environment comprises a processing group with a plurality of members, and wherein the means for detecting comprises means for detecting a failure of at least one of the plurality of members.

12. The system of claim 11, wherein the means for recovering comprises means for synchronizing messages regarding the one or more transactions among surviving members of the processing group.

13. The system of claim 12, wherein the means for recovering further comprises means for committing the one or more transactions.

14. The system of claim 11, wherein at least one member of the processing group survives the failure, and wherein the means for recovering comprises means for electing a coordinator from among the at least one surviving member.

15. The system of claim 14, wherein the means for recovering further comprises means for receiving by the coordinator an indication of the one or more transactions from the other surviving members.

16. The system of claim 15, wherein the means for recovering further comprises means for receiving by the coordinator any commit protocol messages for the one or more transactions the coordinator does not already have.

17. The system of claim 16, wherein the means for recovering further comprises means for the coordinator to initiate the commit protocol for the one or more transactions.

18. The at least one program storage device of claim 3, wherein the shared nothing distributed computing environment comprises a processing group with a plurality of members, and wherein the detecting comprises detecting a failure of at least one of the plurality of members.

19. The at least one program storage device of claim 18, wherein the recovering comprises synchronizing messages regarding the one or more transactions among surviving members of the processing group.

20. The at least one program storage device of claim 19, wherein the recovering further comprises committing the one or more transactions.

21. The at least one program storage device of claim 18, wherein at least one member of the processing group survives the failure, and wherein the recovering comprises electing a coordinator from among the at least one surviving member.

22. The at least one program storage device of claim 21, wherein the recovering further comprises receiving by the coordinator a list of the one or more transactions from the other surviving members.

23. The at least one program storage device of claim 22, wherein the recovering further comprises receiving by the coordinator any commit protocol messages for the one or more transactions the coordinator does not already have.

24. The at least one program storage device of claim 23, wherein the coordinator initiates the commit protocol for the one or more transactions.

25. The method of claim 1, wherein the shared nothing distributed computing environment comprises a distributed synchronous transaction system, and wherein the method comprises a failure recovery method for the distributed synchronous transaction system.

* * * * *